

# A model of dynamic rewiring and knowledge exchange in R&D networks

Mario V. Tomasello<sup>\*,1</sup>, Claudio J. Tessone<sup>1,2</sup> and Frank Schweitzer<sup>1</sup>

<sup>\*</sup>Corresponding author; E-mail: mtomasello@ethz.ch

<sup>1</sup>Chair of Systems Design, ETH Zurich, Department of Management, Technology and Economics, Weinbergstrasse 56/58, CH-8092 Zurich, Switzerland

<sup>2</sup>URPP Social Networks, Department of Business Administration, Universität Zürich, Andreasstrasse 15, CH-8050 Zurich, Switzerland

Received: 7 May 2015

Revised: 18 May 2016

Accepted: 19 May 2016

Published: 17 June 2016

DOI: 10.1142/S0219525916500041

## Abstract

This paper investigates the process of knowledge exchange in inter-firm Research and Development (R&D) alliances by means of an agent-based model. Extant research has pointed out that firms select alliance partners considering both network-related and network-unrelated features (e.g., social capital versus complementary knowledge stocks). In our agent-based model, firms are located in a metric knowledge space. The interaction rules incorporate an exploration phase and a knowledge transfer phase, during which firms search for a new partner and then evaluate whether they can establish an alliance to exchange their knowledge stocks. The model parameters determining the overall system properties are the rate at which alliances form and dissolve and the agents' interaction radius. Next, we define a novel indicator of performance, based on the distance traveled by the firms in the knowledge space. Remarkably, we find that – depending on the alliance formation rate and the interaction radius – firms tend to cluster around one or more attractors in the knowledge space, whose position is an emergent property of the system. And, more importantly, we find that there exists an inverted U-shaped dependence of the network performance on both model parameters.

*Keywords:* Complex network; R&D alliance; knowledge exchange; agent-based model; technological trajectory.

## 1 Introduction

The number of observed inter-organizational Research and Development (R&D) alliances has grown until the 1990s, especially in industrial sectors like IT, Pharmaceuticals and other high-technology ones [1, 21]. This has stimulated research

in different domains, for instance the mechanisms behind the formation of R&D alliances [33], the complex networks they origin [36, 42], and the way they can be described, modeled and their evolution forecasted [25, 14]. A number of theoretical works have shown that, among many other reasons, there exist three main motives for firms to engage in alliances. First, they can gain access to different assets more quickly than they could do *in-house* [29, 9]. Second, alliances foster the exchange of knowledge between firms: by joining their technological resources, firms can actually enlarge their knowledge basis more than they could do individually [5, 30, 35]. Third, firms can share the costs and risks of a project, especially when this is expensive or with uncertain outcome [22]. All of these aspects – even when a knowledge transfer is not directly involved – result in a learning process by the firms [2], of tacit or explicit knowledge, thus making R&D alliances an important element in many firms’ strategy.

In this work we investigate such a learning process, which we model as a knowledge exchange occurring after the establishment of an alliance between two agents. Our agents can change partners and rewire their links as well, thus introducing complex mutual feedbacks between the network structure and their intrinsic characteristics (i.e. their knowledge basis).

Our agent-based model follows an existing stream of literature in the direction of bounded confidence and continuous opinion dynamics models [3, 10, 11, 24, 19], especially applied to innovation networks [13, 4]. In the wake of this previous work, we assume that the collaborating nodes are endowed with an evolving knowledge basis, that affects alliances and – in its turn – is affected by them. However, differently from the studies that have been done so far, our model does not focus on the formation of consensus clusters (see [3, 38] in the case of social systems, or [12] for technology islands). Also, our work differs from previous studies [15, 39] that are focused on strategic decisions made by firms and the effects that these have on the innovation incentives for the involved parties. We rather focus on the dynamics that leads the system to the observed final state, with emphasis on the exploration of the knowledge space by the collaborating agents. We then investigate the existence of an optimal network dynamics that maximizes such a knowledge space exploration.

With respect to R&D networks, it has been shown [36] that – despite long-term simultaneous fluctuations [41] – different industrial sectors exhibit different characteristics in their alliance activity (size and density of the corresponding inter-firm network, heterogeneity of degree distributions, other sophisticated topological network properties and so on). Part of these observed differences have been explained with the so-called “technological regime” of the sector [36]. A technological regime is defined [31] as the pattern of behaviors and common practices in an industrial sector, that are influenced by factors such as technological dynamism, technological uncertainty or separability of innovation activities. In the literature, two technological regimes have originally been detected [45]: an *entrepreneurial regime*, where R&D activities are mainly carried out by new innovative firms, and a *routinized regime*, where innovation is mainly done by incumbent firms. These two extremes are often referred to as *tacit knowledge regime* and *explicit knowledge regime*, respectively, because firms in the network tend to interact with similar or with diverse firms (in terms of knowledge basis), in the respective cases. However, this distinction has been extended over the years, leading to the identification of several classes of technological regimes, spanning between the two aforementioned extremes.

To the best of our knowledge, there is only little research about the influence of technological regimes on the formation of alliances, from a modeling point of view. The present study contributes indeed to this discussion, by developing an agent-based model that reproduces the knowledge exchange process occurring during R&D alliances. Besides, we define here a novel indicator of network performance, based on the exploration of a knowledge space by the agents. In this way, our model is able to capture the existence of an optimal rate of alliance rewiring, as well as its dependence on the underlying technological regime.

## 2 Model foundations

The microscopic rules of our agent-based model are inspired by a number of stylized facts, as well as theoretical speculations, in network evolution studies, opinion dynamics models, R&D and collaboration networks. Below, we provide a brief description of every building block that we employ in the development of our model.

**Monogamous network approximation.** We model the formation of R&D alliances between companies as a *monogamous network*, i.e. a network in which every agent is linked to only one other agent at every time step [44, 40]. Inter-organizational networks are indeed proven to have low density, i.e. only a small fraction of all potential collaborations between companies are actually realized. The density of R&D networks ranges from 0.1% to 1% for all industrial sectors, as shown in a previous study [41]. However, some high-technology industries, such as Pharmaceuticals or Computers, although having low density, show high clustering and hierarchical structures, which are of fundamental importance for the dynamics of knowledge diffusion. The “hubs” of these industrial sectors can actually have more than a hundred partners at the same time, with which they collaborate on different projects [33, 23].

Despite this empirical fact, we still use here the monogamous network as a modeling tool. In order to have a more realistic picture, it should be noted that – even though the agents have only one link at every time step – they are allowed to change their partners in the following steps and can actually collaborate with many firms in a small time window. Therefore, we propose as a possible extension to aggregate many network snapshots over time, similar to a previous theoretical study [4] on R&D networks; however, we deem this investigation to be beyond the scope of the present model and leave it for future research.

**Position of companies.** In the knowledge-based view of the firm, every company is endowed with a knowledge basis that uniquely identifies its resources and its capabilities. We assume that a firm is represented by an agent in our modeling framework, and associate it with a vector of  $D$  components, each of which represents its share of knowledge in a given area. Furthermore, we directly associate these vectors to a metric *knowledge space* in which the collaborations occur: every firm occupies a point in this  $D$ –dimensional space, whose coordinates are given by its knowledge vector. Such an approach is similar to a more general model [3], proposed in the broader context of social influence. The concept of a metric knowledge space has already been used in one dimension [19],

and in two dimensions [12, 4]. We generalize the dimensionality of the space to  $D$ .

The coordinates of every node can be thought of as the ratios of the corresponding firm’s expertise along each of the  $D$  dimensions of the space. In order to have an empirical representation of these ratios, following an existing study [43], one can think of the different technological classes of which the International Patenting Classification (IPC) scheme is composed. Just to give an example, the real IPC is divided into eight main categories, spanning from “human necessities” to “electricity”.

Assuming that the classification scheme for our firms consists of  $D'$  categories, then the  $D$  values (with  $D = D' - 1$ ) would be the fractions of patents in each category to the total number of patents issued by the firm at hand. It is important to note that such values are ratios, and not absolute measures of knowledge; therefore, there are no *better* positions than others in the knowledge space that we utilize, but only *different* positions, between which we can easily compute an appropriate measure of similarity. These  $D$  ratios are free to vary independently of each other in the interval  $[0, 1]$ ; the remaining  $D'$ -th, or  $(D+1)$ -th, knowledge component can be inferred from the main  $D$  values through the bounding condition that the  $D'$  values have to sum up to 1.

**Alliance formation.** In our monogamous network, all nodes are linked in pairs at every time step. We assume that two pairs of allied nodes mutually rewire their links at every time step with a given probability, and the new formed links are active if the Euclidean distance between the new partners is smaller than a threshold value. Such a proximity condition models some existing theoretical arguments [7, 6], highlighting that an interaction between two companies is profitable only if their *absorptive capacity* is large enough or – in other words – their knowledge distance is small enough. However, further studies [28, 18] have shown that there exists an inverted U-shaped relationship between the profitability of an alliance and the knowledge distance of two companies. This means that, partners with a *too small* knowledge distance (in other words, a *too high* similarity) do not have any reason to establish an alliance. Even though the selection strategy does not include such a curvilinear dependence, our agent-based model is able to capture this stylized fact, because we assume that the learning speed of the two agents (see Section 3 for more details) decreases with their knowledge distance, i.e. their learning potential. Therefore, partners which are already similar in terms of knowledge bases do not significantly contribute to the increase of the relevant performance indicator in our model.

The choice of the Euclidean metric to compute this distance is quite realistic, even if it implies extensive information about the companies’ mutual position in the knowledge space. Indeed, obtaining detailed information about a company, its patent production, its scientific production and its activities in general is nowadays not only feasible – thanks to the Internet – but actually done by most firms willing to engage in an alliance [1, 37, 4].

The threshold value for the alliance establishment is supposed to model the technological regime that characterizes the collaboration network under examination. A large interaction threshold means that the agents can establish active collaborations even with agents located far away in the knowledge space; this corresponds to an explicit knowledge regime, typical of a mature industry, where

innovation is more routinized and mainly carried out by large incumbent firms, which have easy access to both similar and different firms (in terms of knowledge basis). A small interaction threshold means that the agents can establish active collaborations only with agents located close in the knowledge space; this corresponds to a tacit knowledge regime, typical of a young industrial sector, where innovation is mostly carried out by small new-entrant firms, which have easy access to similar others in terms of knowledge basis.

**Partner selection.** The dynamics of alliance formation in the present model is assumed to be *semi-random*, meaning that the rewiring of links between nodes occurs randomly and independently of the position of the nodes themselves in the knowledge space: we call this an *exploration phase*. However, a link between two nodes is *active* only if they are close enough in the knowledge space: if this happens, a so-called *knowledge transfer phase* begins. The rewiring mechanism does not intend to be a close representation of what happens in reality. It rather has the function of modeling the volatility of R&D alliances, capturing the characteristic time scale at which firms decide to engage in a new alliance. The second focal aspect that we want to model – namely the formation of alliances at the right knowledge distance – is instead fully captured by the threshold value for the potential partner’s knowledge similarity.

**Approaching in the knowledge space.** Once a link has been established, we assume that a knowledge exchange between the partners takes place, causing their knowledge bases to become more similar and making them approach in the knowledge space. This assumption is in line with the conceptualization of R&D alliances as a means to exchange technological knowledge among firms [30, 32, 17] and has already been used in a number of agent based models [34, 16, 8]. Besides, we argue that even when a knowledge transfer is not directly involved, for instance in a marketing or a cost-reducing alliance, an exchange of tacit knowledge can still take place between the two interested companies.

In our agent-based model, the speed at which the agents approach each other – or, in other words, the rate at which they mutually learn from each other – is governed by one parameter (further explanations follow below). The closer the two agents are, the smaller their approaching speed becomes, as they are depleting the potential for mutual learning. In addition, it should be noted that our work studies a scenario in which the knowledge spillovers occurring in a R&D alliance cause the partners to exchange knowledge along every dimension, not limiting the knowledge transfer to a single area of expertise – i.e. one of the  $D$  dimensions in the knowledge space. Practically, this means that in every time step of our computer simulations the  $D$  knowledge ratios of every pair of allied firms modify their values and become more similar; this approach is similar to a previous model [4].

**Exploration of the knowledge space.** Finally, we want to study the performance of the whole collaboration network as a function of the relevant model parameters. The indicator we propose to measure such a performance takes into account the global knowledge exploration of the system, i.e. it quantifies the distance traveled by all agents during the evolution of our simulated R&D network. In our model, we consider that the knowledge exploration itself is

represented by the motion in the space, which is fully captured by this indicator. The underlying assumption is that the exploration of as many locations as possible is beneficial for the R&D collaboration network, in that it allows the agents to come in contact with many technological opportunities, potentially leading to more frequent innovations [12]. Testing our model by means of computer simulations, we find that the rewiring of links and the mutual knowledge exchanges over time eventually lead the whole system to a steady state through a peculiar dynamics. The model and its results are presented in detail in the next two Sections.

### 3 The model

Starting from the evidence and the arguments presented in the previous Section, we now present the implementation of the agent-based model. We consider a network composed of  $N$  nodes, each representing an agent – in the particular case of R&D networks, a firm – performing collaboration activities in a knowledge space. The model is implemented by means of computer simulations, consisting of a sequence of discrete time steps of length  $dt$ . The microscopic interaction rules are described below.

#### 3.1 Exploration phase

Every node  $i$  is located in a metric space (henceforth, the knowledge space); this point has coordinates  $\mathbf{x}_i$ , identified by a vector of  $D$  real numbers ranging from 0 to 1. As already explained, the coordinates of every node can be thought of as the ratios of the corresponding firm’s expertise along each of the  $D$  dimensions of the knowledge space. At the initial stage of every simulation, all the nodes’ positions are drawn from a uniform distribution in the interval  $[0, 1]$ .

$$\mathbf{x}_i \equiv (x_{i1}, x_{i2}, \dots, x_{iD}) \quad i = 1, \dots, N. \quad (1)$$

All nodes in our R&D network have the possibility to change their partner, thus generating a dynamic network topology. We model this by means of a link rewiring mechanism. The time steps in our computer simulations have a duration equal to  $dt$ ; in each time step, two pairs of connected firms are randomly chosen and, with a rate  $\lambda$ , they rewire their links. We call this process “exploration phase”, and depict it in Fig. 1. Let us assume that the nodes  $i$  and  $j$  and the nodes  $i'$  and  $j'$  constitute the two linked pairs chosen at time  $t$ . With probability  $\lambda dt$ , they mutually exchange their partners, and at time  $t + dt$  the nodes  $i$  and  $i'$  and the nodes  $j$  and  $j'$  will form the new linked pairs. With probability  $1 - \lambda dt$ , instead, nothing happens and at time  $t + dt$  the nodes  $i$  and  $j$  and the nodes  $i'$  and  $j'$  will still be respectively linked.

Such a random search for partners in the exploration phase has the function to model the volatility of R&D alliances, capturing the characteristic time scale at which an agent decides to engage in a new collaboration. The rate  $\lambda$  can be indeed thought of as the inverse of the characteristic time elapsed before a firm takes part in a new alliance. Even though the potential partner is selected at random, the R&D alliance will be actually “active” only if the partner fulfills a certain proximity condition in the knowledge space, as we will explain below.




Figure 1: Schematization of a link rewiring between two pairs of connected nodes. At time  $t$ , the nodes  $i$  and  $j$  and the nodes  $i'$  and  $j'$  are linked in pairs. These two couples of nodes are selected and, with probability  $\lambda dt$ , they switch links: at time  $t + dt$  the nodes  $i$  and  $i'$  and the nodes  $j$  and  $j'$  are the new linked pairs. Obviously, with probability  $1 - \lambda dt$ , no rewiring happens.

Therefore, such an exploration is not fully arbitrary, and leads to the establishment of an actual collaboration only under specific conditions. It is worth mentioning that the results of our simulations remain qualitatively unchanged if we use any different random link creation process, as demonstrated in a previous study [40].

### 3.2 Knowledge transfer phase

The whole linking and rewiring process in our model occurs independently of the node knowledge positions, but their distance in the knowledge space has a determinant effect on the subsequent network dynamics. Indeed, one of the key ingredients of our model is the existence of an optimal absorptive capacity for a profitable R&D alliance between two firms. We assume that a link is *active* if the corresponding pair exhibits a knowledge distance smaller than a given threshold value. If this proximity condition is not fulfilled, even though the corresponding nodes are connected, their link is considered to be *inactive*, causing no effect at all on the system. The proximity condition is evaluated for every pair of linked nodes  $i$  and  $i'$  as follows:

$$|\mathbf{x}_i(t) - \mathbf{x}_{i'}(t)| < \varepsilon\sqrt{D} \quad (2)$$

where we employ the Euclidean distance  $|\cdot|$ , consistently with the assumption of evaluating the diversity of each firm's knowledge portfolio in all dimensions.  $\sqrt{D}$  is the maximum possible distance between two points in a  $D$ -dimensional Euclidean space. The parameter  $\varepsilon$ , ranging from 0 to 1, is the threshold interaction radius inside which nodes are able to interact and collaborate profitably. Only links whose corresponding nodes fulfill this proximity condition are considered to be *active*. Such an interaction radius can be associated with the knowledge regime characterizing the collaboration network under examination. A large  $\varepsilon$  means that the firms can potentially see and explore a large portion of the knowledge space, being the knowledge highly *codified*. A small  $\varepsilon$  represents instead a regime of *tacit* knowledge, where firms are able to establish alliances only if their technological positions are already close.

We assume that an R&D alliance causes the two involved firms to pool their resources and their knowledge basis, thus approaching along every dimension in the knowledge space. Thanks to *knowledge spillovers*, both firms will acquire common practices or a shared jargon, not limiting the knowledge transfer to that specific R&D project that they have in common, as previously discussed.<sup>1</sup> If  $i$  is an agent and  $i'$  is its unique partner in the collaboration network at time  $t$ , both will move towards each other by identical paths in the knowledge space, provided that the proximity condition expressed in Eq. 2 holds. The model dynamics equation is the following:

$$\dot{\mathbf{x}}_i(t) = \mu [\mathbf{x}_{i'}(t) - \mathbf{x}_i(t)], \quad \text{if } |\mathbf{x}_{i'}(t) - \mathbf{x}_i(t)| < \varepsilon\sqrt{D} \quad (3)$$

where  $\mu$  is defined as the *learning rate* of the agents. This parameter is constant over time and for all nodes in the collaboration network, and can be thought of as the propensity of the agents to exchange knowledge with their partners, thus making their knowledge bases more similar over time. It should be noted that the parameter  $\mu$  is a rate, not a speed; the actual speed at which the corresponding nodes move in the knowledge space is given by the product of the rate  $\mu$  and their distance: therefore, the farther they are in the knowledge space, the faster they approach. When their distance decreases, so does the potential for new learning from the collaboration, and the approaching speed drops consequently. This interpretation is clear in Eq. 4, which represents the way we implement the model in computer simulations with discrete time steps of length  $dt$ . The evolution of every agent's position  $\mathbf{x}_i$  can be expressed as:

$$\mathbf{x}_i(t + dt) = \mathbf{x}_i(t) + \mu dt [\mathbf{x}_{i'}(t) - \mathbf{x}_i(t)] \quad (4)$$

We depict such knowledge exchange mechanism in Fig. 2. The nomenclature and the meaning of all the model parameters we introduced in this Section are summarized in Table 1.

Parameter	Meaning	Type of parameter
$N$	Number of agents (system size)	Static
$D$	Dimensionality of the metric knowledge space	Static
$\varepsilon$	Agents' interaction radius (knowledge regime)	Static
$\lambda$	Link rewiring rate	Network dynamics
$\mu$	Approaching rate in the knowledge space	Network dynamics

Table 1: Model parameters and their description. The “static” parameters are associated with the system structural features, while the “network dynamics” parameters define the characteristic speed at which the system evolves.

## 4 Results

In order to test our model, we have performed extensive computer simulations by applying the dynamics presented in Section 3 and varying the values of the relevant parameters. In particular, we vary the size  $N$  of the network from 10

---

<sup>1</sup>However, we have also tested a scenario in which two allied firms exchange knowledge only in one dimension, thus moving in only one dimension of the knowledge space as well. The results remain qualitatively unchanged.






Figure 2: Schematization of the knowledge exchange process in a bi-dimensional space ( $D = 2$ ). At time  $t$ , the agents  $i$  and  $i'$  are linked and their distance  $|\mathbf{x}_{i'}(t) - \mathbf{x}_i(t)|$  is smaller than  $\varepsilon\sqrt{D}$ ; consequently, at time  $t + dt$ , their positions  $\mathbf{x}_i(t + dt)$  and  $\mathbf{x}_{i'}(t + dt)$  will be closer in the knowledge space. The picture includes other pairs of connected agents, whose distance is larger than  $\varepsilon\sqrt{D}$ . Therefore, these links are *inactive* (depicted in dashed red lines) and do not originate any motion in the knowledge space.

to 2,000 nodes, the dimensionality  $D$  of the knowledge space from 1 to 50, the interaction threshold radius  $\varepsilon$  from 0 to 1, the learning rate  $\mu$  from  $10^{-3}$  to  $10^3$  and the rewiring rate  $\lambda$  from  $10^{-3}$  to  $10^4$ . All of these parameters are explored in discrete intervals, whose width is appropriately chosen – as we discuss below in more detail. The value of  $dt$  is instead fixed for all our computer simulations to a value of 0.0001. All values are expressed in arbitrary units.

**Main model parameters and their meaning.** We argue that the network evolution is essentially characterized by two driving forces with overall opposite effects. The first one is the formation of active links (i.e. the establishment of profitable alliances or collaborations); this force tends to push agents closer in the knowledge space, given the resulting approaching motion. The second force is the link rewiring (representing the dissolution of old collaborations and the formation of new ones), that stimulates the agents to explore new portions of the knowledge space. This force could result in a faster overlap of every agent's knowledge position, but it could also result – under certain conditions – in preventing the agents from converging to a knowledge attractor, thus keeping them far-between in the knowledge space.

These competing forces are associated with the two model dynamics parameters, respectively the approaching rate  $\mu$  and the link rewiring rate  $\lambda$ . However, it is clear that the relation between these two parameters will substantially affect the emergent properties of the system. What truly affects the resulting dynamics of the network are not the absolute values of the two rates  $\mu$  and  $\lambda$ , but the ratio of the two. Indeed, using a configuration with the same  $\mu$  to  $\lambda$  ratio, but with smaller absolute values, will only lead to a longer computer simulation (i.e.

more discrete time steps are needed), without qualitatively changing the results. Therefore, in the continuation of the current study we present our findings by keeping the value of the learning rate fixed to  $\mu = 1$ , and studying the effect of the dynamics parameter  $\lambda$  only.

The second relevant model parameter on which we focus our attention is the threshold interaction radius  $\varepsilon$ , a static parameter representing the knowledge regime in which the collaborating agents move. We explore a series of values ranging from a totally tacit knowledge regime ( $\varepsilon = 0$ ) to a totally explicit one ( $\varepsilon = 1$ ).

The effect of the dimensionality parameter  $D$  and the network size  $N$  is that of changing the characteristic density of the system, i.e. the number of agents that can be found within the given interaction radius  $\varepsilon$ . Varying their values causes shifts in the trends of the relevant measures that we investigate (network performance, number of knowledge clusters and convergence time), without qualitatively affecting the results. Indeed, the numeric values of the parameters that we present here do not have a specific meaning, and the ranges that we explore in the present study are only aimed at a clear and effective visualization of our findings. For the sake of extendability and future empirical work, all our model parameters have a straightforward and natural match with real quantities, including the parameters whose dependencies are not directly studied here, such as the size  $N$  of the system, its dimensionality  $D$ , or the learning rate  $\mu$ . Should these parameters be matched with their empirical counterparts, all their values would assume a real, directly intelligible meaning.

One last remark has to be made about the number of agents  $N$  in the system. In empirical collaboration networks, such a number is obviously dynamic, and not static. However, the incorporation of a dynamic system size in our agent-based model would deeply modify it, and shift the focus away from the investigation of the agents' knowledge exploration. Such an extension could indeed constitute a second, distinct study, and we leave it for future research.<sup>2</sup>

In the continuation of the present study, we select a network composed of  $N = 200$  nodes and a knowledge space with  $D = 10$  dimensions, to present our results in the most effective way possible.

**Network performance.** The variable that we investigate as indicator of the network performance is the *mean knowledge path*  $\langle K \rangle$  of the collaborating agents. We define the path covered by every agent in the knowledge space  $K_i$  as the sum of all the distances that the agent travels in every time step of the simulation:

$$K_i = \int_{t=0}^{T_{\max}} |\dot{\mathbf{x}}_i(t)| dt \quad (5)$$

where  $T_{\max}$  is the duration of an entire computer simulation. It should be noted that the measure  $|\dot{\mathbf{x}}_i(t)| dt$  is a positive scalar and expresses the actual distance traveled by the agent  $i$ , differently from its net displacement  $\dot{\mathbf{x}}_i(t) dt$ , which is a vectorial quantity. The measure  $K_i$  is then averaged over all the  $N$  network agents to obtain the mean knowledge path  $\langle K \rangle = N^{-1} \cdot \sum_i K_i$ . We hypothesize

---

<sup>2</sup> Indeed, there exists a work [43] which attempts to study a similar issue of knowledge exchange in a system where the time scales of the agents' interactions and their entry/exit in the network is not decoupled; a dynamic number of agents can partially be captured by that model through a quantity called "activity". Such a study could be extended and improved with the findings deriving from the present model.

that this measure can provide a meaningful indication of the macroscopic system performance, because – as already discussed in Section 2 with respect to the microscopic level – firms are proven to innovate more when they come in contact with more technological opportunities. Therefore, we assume that a higher value of  $\langle K \rangle$ , i.e. a higher distance explored in the knowledge space, corresponds to a higher network performance. We argue that the same reasoning can be as well extended to other types of collaborations that involve learning and/or knowledge exchange processes.

We present the results in Fig. 3, for a representative network of  $N = 200$  agents moving in a knowledge space with  $D = 10$  dimensions. As already mentioned, the parameter  $\mu$  is fixed to 1, and we study the dependence of  $\langle K \rangle$  on the dynamics parameter  $\lambda$  and the static parameter  $\varepsilon$ .




Figure 3: Mean knowledge path  $\langle K \rangle$  (displayed by means of both the z-elevation and the color scale), as a function of the rewiring rate  $\lambda$  and the interaction radius  $\varepsilon$ . The R&D network under examination has  $N = 200$  nodes and learning rate  $\mu = 1$ , in a 10-dimensional knowledge space. We generate 1000 simulations for each parameter set and then average the results.

We find that the mean knowledge path  $\langle K \rangle$  exhibits a peak in correspondence of specific values for both the rewiring rate  $\lambda$  and the interaction radius  $\varepsilon$ . In the case that we present in Fig. 3, these values are  $\lambda \simeq 10$  and  $\varepsilon \simeq 0.25$ , located in the red area of the plot.

Taking a closer look at the network performance, we find that  $\langle K \rangle$  shows a monotonic growing trend as a function of  $\lambda$ , when the interaction radius  $\varepsilon$  is lower than a certain value  $\varepsilon^*$  (in our example,  $\varepsilon^* \simeq 0.25$ ). When fixing the interaction radius to larger values  $\varepsilon \geq \varepsilon^*$ , we do instead find that  $\langle K \rangle$  exhibits a peak as a function of  $\lambda$ . This means that, as the knowledge regime becomes more explicit, and the agents are allowed to form active collaborations even with distant partners in terms of knowledge basis, there exists a rewiring rate

maximizing the distance actually explored by the agents in the knowledge space.

The behavior of the mean knowledge path  $\langle K \rangle$  can also be described as a function of the interaction radius  $\varepsilon$ , while keeping the rewiring rate  $\lambda$  fixed. We find that  $\langle K \rangle$  grows with  $\varepsilon$  to a saturation level (when  $\varepsilon > 0.5$ ), if the rewiring rate is small ( $\lambda < 1$ , for the case under study). If we fix the rewiring rate  $\lambda$  to a value larger than 1, we find instead that  $\langle K \rangle$  increases to a peak, in correspondence to  $\varepsilon \simeq 0.35$ , and then decreases again to stabilize for  $\varepsilon > 0.5$ . This means that, when the characteristic alliance rewiring rate of the network is greater than the characteristic learning rate of the agents – i.e. when the search for new alliance partners is predominant over the learning mechanism – there exists a specific threshold interaction radius (corresponding to a moderately explicit knowledge regime) maximizing the distance covered by the agents in the knowledge space.

This finding can be interpreted in terms of emergence of knowledge attractors in the system (please refer to the paragraph below for more detailed explanations). In a regime dominated by the exploratory search for new alliance partners ( $\lambda > 1$ ), increasing the interaction radius  $\varepsilon$  is beneficial for the system, because the agents can form alliances with more distant partners, and therefore travel longer distances in the knowledge space. However, a too high interaction radius – together with  $\lambda$  values greater than 1 – causes the emergence of just one central attractor, toward which the agents converge quickly, without traveling too much distance in the knowledge space. This happens because the agents have the chance to interact with all others, thus quickly uniforming their knowledge bases.

On the other hand, when looking at the rewiring rate rather than the interaction radius, the explanation of the inverted U-shaped network performance is dynamical instead of spatial. Here, a high rewiring rate is beneficial because it brings the agents in rapid contact with other agents situated far away in the knowledge space, increasing – on average – the distance that they travel. However, a too high rewiring rate (combined with an intermediate value of  $\varepsilon$ ) results in the emergence of several attractors in distinct locations of the knowledge space, as we report in detail in the next paragraph. This causes each of the agents to converge to one of the many attractors, rather than the unique central attractor, which would instead emerge in case of a lower rewiring rate. Eventually, this translates into a lower distance traveled by the agents, if compared to a medium-rewiring-rate scenario, where all agents would globally travel a longer distance to reach the central attractor.

Whether the distance traveled in the knowledge space is a better performance indicator than the number of emerging knowledge attractors is still an open question, which probably requires a case-by-case discussion depending on the system under examination. In any case, the analysis of the mean knowledge path in the system cannot be decoupled from the analysis of the knowledge attractors emerging in the system. In the next paragraph, we examine in more detail how our agent-based model can explain their formation and evolution.

**Knowledge clusters and attractors.** We here investigate a second emerging property of the system, namely the *number of knowledge clusters* appearing in the network at the end of every model run. We define a knowledge cluster as a group of nodes whose mutual distances are smaller than  $\varepsilon$ . Moreover,

the distance between every node in that cluster and every node outside of that cluster has to be larger than  $\varepsilon$ , meaning that all the agents in the cluster will asymptotically converge to one point in the knowledge space, and no further inclusion of any other agent in the cluster is possible. We call such a point a knowledge attractor, or simply an attractor.

It is clear that the maximum possible value of knowledge clusters equals the number of nodes  $N$ ; we expect to observe such a value in correspondence with a low value of the interaction radius  $\varepsilon$ , when the agents are virtually unable to establish active links. Likewise, the minimum possible number of knowledge clusters equals 1; we expect to observe such a value in correspondence with high values for the interaction radius  $\varepsilon$ , when most established collaborations are active, thus facilitating the convergence of all agents toward one knowledge attractor. Similarly to the mean knowledge path, we present our results in Fig. 4, for a network of  $N = 200$  agents in a knowledge space with  $D = 10$  dimensions;  $\mu$  is fixed to 1.




Figure 4: Number of knowledge clusters as a function of the rewiring rate  $\lambda$ , for a set of representative values of the interaction radius  $\varepsilon$ . The network under examination has  $N = 200$  nodes and learning rate  $\mu = 1$ , in a 10-dimensional knowledge space. We generate 1000 simulations for each parameter set and then average the results.

We find that the number of clusters generally increases by decreasing the interaction radius  $\varepsilon$ . As expected, an extreme case occurs for  $\varepsilon = 0$  (completely tacit knowledge regime, where any interaction is by definition impossible), in which we have as many clusters as agents – independently of the rewiring rate  $\lambda$ . The other extreme case occurs for  $\varepsilon \geq 0.5$  (highly explicit knowledge regime), in which all the nodes interact between each other converging in only one cluster – again, independently of  $\lambda$ .

Noteworthy, for intermediate values of  $\varepsilon$ , we observe an interesting dependence of the number of knowledge clusters on the rewiring rate  $\lambda$ . When  $\lambda$  is low (i.e. comparable in magnitude with  $\mu$ ), we find the existence of one or very few knowledge clusters, because the overall effect of such a slow rewiring rate is that all nodes tend to get closer in the knowledge space before the corresponding links are cut and rewired. As a result, all nodes are eventually part of the same

knowledge cluster. From the visual examples in Fig. 5 (a) and (b), we can observe that such clusters are dispersed in the knowledge space, and the presence of a central attractor is not visually detectable, even though all the agents are in principle within interaction distance. What happens, in fact, is that every pair of agents converges to the midpoint of the segment connecting them; the system then “freezes” in this configuration, being the rewiring rate too low to allow for new collaborations and new explorations within a meaningful time frame.

It should be noted that, if the computer simulations were allowed to last longer, the agents could in principle converge to the central attractor above mentioned, thus making it better visible in the knowledge space. However, this would result only in a different visual outcome and would change neither the number of detected knowledge clusters (all the agents are included in one giant cluster anyway; see definition above), nor the way the agents travel in the space (i.e. converging to the midpoint of the connecting segments and then freezing there for most of the time), thus not affecting our findings about the network performance.

When the value of  $\lambda$  increases, instead, we observe the formation of a higher number of knowledge clusters. These clusters are well delimited in the knowledge space and, as we show in the examples of Fig. 5 (c) and (d), the presence of attractors is visually evident. Such a non-trivial effect derives from the fact that the nodes cut their links and form new ones before the approaching mechanism with the previous partner is complete, thus traveling with a peculiar meandering trajectory, also clearly visible in Fig. 5 (c) and (d).

It is interesting to note how such peculiar trajectories result both in a longer traveled distance in the knowledge space and in the emergence of several attractors, occupying different regions of the space. However, when  $\lambda$  increases above a given threshold – in accordance with our findings on the network performance – the longer meandering trajectories are no longer able to compensate for the shorter distances globally traveled by the agents. The emergence of several distinct attractors, indeed, causes the agents to travel a shorter distance, on average, before converging to one of them. For the case examined in Fig. 3, this results in a decreasing network performance when the rewiring rate is higher than  $\sim 10$ , interestingly the value in correspondence of which the number of clusters explodes (see Fig. 4).

**Convergence time.** We find that the network dynamics generated by the model eventually converges to a steady state, in which all the agents occupy one or more fixed positions, and no further collaborations, nor motion in the knowledge space, are possible. In other words, such a steady state represents a configuration in which the collaborating agents have depleted all the potential for new knowledge exchange.

We define a convergence criterion based on the agents’ motion in the knowledge space, and assume that the steady state is reached if the total knowledge path traveled by all the agents in the last time step is smaller than the 0.5% of the knowledge path covered by the agents in the last 500 time steps. Indeed, all of the network measures described above are computed only after the steady state is reached. In Fig. 6, we show the trend of the convergence time as a function of  $\lambda$  and  $\varepsilon$ , for the same representative network that we have studied before.

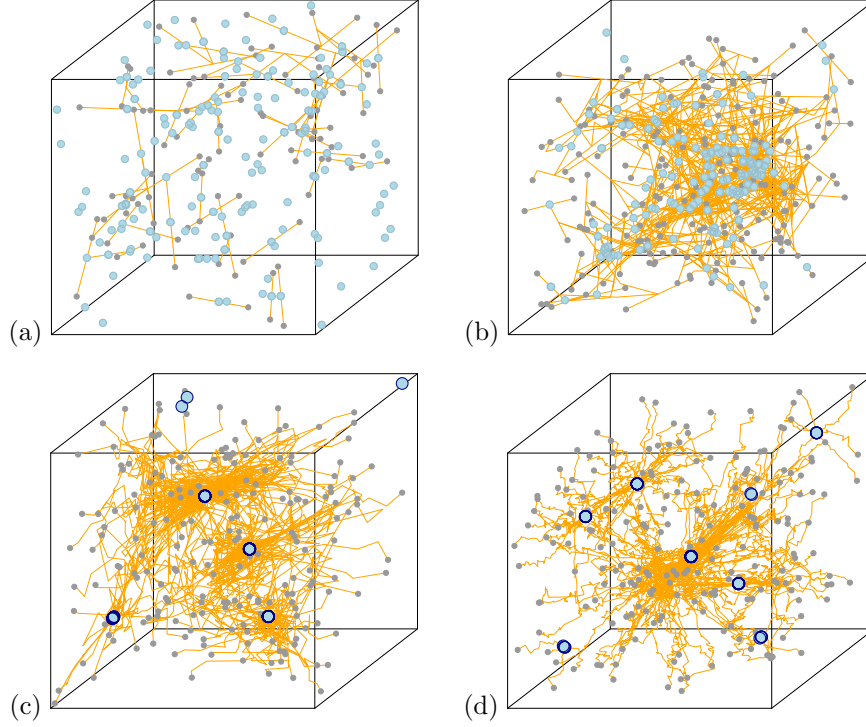


Figure 5: Knowledge trajectories for a network with  $N = 200$  nodes and learning rate  $\mu = 1$ . For the sake of visualization, here we use a knowledge space with  $D = 3$  dimensions, easily representable as a cube. The initial positions of the nodes are depicted with gray dots, their trajectories with orange lines, and their final positions with blue dots. If the final position of an agents corresponds to a knowledge attractor, we depict this with a circled, dark blue dot. We keep the threshold interaction radius constant to  $\varepsilon = 0.2$ , and show four cases corresponding to rewiring rate  $\lambda$  equal to: (a) 1, (b) 10, (c)  $10^2$ , (d)  $10^3$ . In the cases (a) and (b), the agents are “frozen” and dispersed in the knowledge space, and they all belong to a unique, giant knowledge cluster. In the cases (c) and (d), the faster rewiring rate allows for the emergence of several distinct knowledge attractors, to which the agents converge through longer, meandering trajectories.

On the one hand, we find that all the relevant parameter configurations reach a steady state before the computer simulation ends. Indeed, the parameter combinations that are not able to reach a steady state before the end of the simulation (those with  $\varepsilon < 0.15$  or generally low  $\lambda$ ) are the ones generating the lowest values of mean knowledge path, for the reasons we previously discussed. Therefore, we forcibly stop all computer simulations after 20,000 time steps, affecting only a small fraction of the parameter space and not influencing our results. It should be noted that we aim at developing a model that has the potential to be validated against real data. Because of this, we are interested in configurations which can be considered as a pseudo steady state within a




Figure 6: Convergence time as a function of the rewiring rate  $\lambda$ , for a set of representative values of the interaction radius  $\varepsilon$ . The network under examination has  $N = 200$  nodes and learning rate  $\mu = 1$ , in a 10-dimensional knowledge space. We generate 1000 simulations for each parameter set and then average the results.

meaningful time frame; this means a few orders of magnitudes ( $\sim 4$ , in our case) longer than the characteristic interaction time. This is consistent with most of the empirical datasets on R&D or collaboration networks, whose typical observation length (a few decades) is around 4 orders of magnitude larger than their granularity (1 day).<sup>3</sup> Therefore, when true convergence to a stationary state requires letting the system evolve for a much higher number of time steps, this could be deemed unrealistic.

On the other hand, we find an unexpected trend of the convergence time as a function of  $\lambda$  for some parameters combinations. One would expect that the convergence time decreases proportionally to  $1/\lambda$ , being the inverse of the rewiring rate a measure of the characteristic time of the system for a complete interaction between all agents. More precisely, considering that  $\lambda$  is the characteristic rate for one interaction between two agents in the system, its inverse  $1/\lambda$  is the time needed, on average, by each agent to interact with all other agents in the system. This means that, if one neglects all the network effects and the complex interdependencies of the model, the characteristic convergence time of the system would be proportional to  $1/\lambda$ .

However, we observe such a trend only for the extreme cases of highly explicit knowledge regimes, corresponding to  $\varepsilon \geq 0.5$ , where the network effects are minimum and a complete interaction between all agents does indeed take place, because of the very large interaction radius. Instead, we find a peculiar trend of the convergence time as a function of  $\lambda$  for all the other values of  $\varepsilon$ , showing

<sup>3</sup> In our model, the quantity  $\lambda \cdot dt$  represents, at its core, a simplification of a Poisson process. In a hypothetical future work, when validating the model on empirical data, one should pay careful attention to the utilized numeric values. For instance, if one assumes that a simulated time step equals one day in reality, the values of both  $dt$  and  $\lambda$  have to be chosen in such a way that the number of simulated rewirings is comparable with the number of empirically observed inter-firm alliances in the given observation period, to allow for a matching of the granularity of the events.



plateaux for high values of  $\lambda$ . This means that the complex network dynamics, in the presence of certain approaching and link rewiring rates, can slow down the convergence of the system; as a result, the steady state is reached later than the characteristic time  $1/\lambda$  would suggest.

## 5 Discussion

In this study, we have developed an agent based model of knowledge exchange and dynamic rewiring of R&D alliances. Our novel contribution has been to explicitly model how agents move toward each other in a metric knowledge space. In addition, we have studied the co-evolution and the interdependencies of such a process with a dynamically evolving network structure.

By studying the interactions of a set of agents in a metric knowledge space via computer simulations, we have found that the system follows a peculiar dynamics and reaches a steady state in which the agents cluster around a set of emerging attractors. The model parameters that determine the overall properties of the system are the link rewiring rate of the network and the agents' interaction radius.

**Our findings.** We have defined a knowledge cluster as a group of agents whose mutual distances are smaller than the threshold interaction radius, and whose distance with every node outside the cluster is larger than this radius (meaning that all the agents in the cluster will asymptotically converge to one attractor and no further inclusion of any other agent in the cluster is possible). We have found that the number of knowledge clusters observed at the end of the network evolution decreases by increasing the threshold interaction radius, because the agents are able to collaborate with partners located farther away in the knowledge space, thus converging all together towards one position.

When the knowledge regime is strongly tacit or strongly explicit, the number of knowledge clusters depends only on the interaction radius itself, and not on the alliance rewiring rate. The most interesting case occurs for intermediate knowledge regimes, in which the number of knowledge clusters increases with the rewiring rate. Small rewiring rates lead to the emergence of only one knowledge cluster, which is dispersed in the knowledge space and does not clearly exhibit the presence of a knowledge attractor. Faster alliance rewiring rates allow the emergence of a larger number of knowledge clusters; in this case, the presence of knowledge attractors, around which the firms eventually cluster, is (even visually) clear. In such a regime, the agents travel – on average – longer distances in the knowledge space. However, if the rewiring rate is too high, the effect of having more knowledge attractors in the system is detrimental for the agents, which do eventually explore a shorter distance.

**Our interpretation.** The underlying assumption of our agent-based model is that the exploration of as many locations as possible is beneficial for the entire collaboration network. For this reason, we consider the distance explored by the agents in the knowledge space as a *performance indicator* of the network evolution. We have found that there exists an inverted U-shaped dependence of such an indicator on both the alliance rewiring rate and the interaction radius.

In particular, if we focus on the dependence of the performance on the rewiring rate, as already mentioned, we find that there exists a specific value of the rewiring rate maximizing the performance. Such a rate exhibits a weak dependence on the interaction radius; namely, it slightly decreases when the radius increases (only for intermediate radius values). This is consistent with some empirical studies [36, 20], that show a varying alliance formation rate across industrial sectors. Similarly, we have found that, given a fixed alliance rewiring rate, there exists a value of the interaction radius maximizing the network performance.

From the point of view of our agent-based model, this happens because in a regime dominated by the exploratory search for new alliance partners, a high rewiring rate allows the agents to cut their links and form new ones before the approaching mechanism with the previous partner is complete, thus traveling with a peculiar “zig zag”, meandering trajectory. This has a beneficial effect on the system’s performance, because such peculiar trajectories result in a longer traveled distance in the knowledge space.

However, when the rewiring rate is too high, the longer meandering trajectories are no longer able to compensate for the shorter distances globally traveled by the agents, due to the emergence of too many attractors (the agents have to travel less, on average, before converging to one of them). As we have already detected, this results in an inverted U-shaped dependence of the network performance on the rewiring rate.

The same dependence of the network performance occurs as a function of the interaction radius. A higher radius allows the agents to form links even with agents located very far away in the knowledge space, causing a higher distance to be traveled, on average. A too high radius, on the other hand, causes the emergence of fewer attractors, because the agents can interact with more potential partners, thus converging faster to one or few attractors. This results in a lower distance globally traveled by the agents to reach the one (or few) attractors, as opposed to the longer distance that they would travel to reach the many different attractors in a scenario with medium interaction radius (and medium/high rewiring rate).

**Conclusions.** In conclusion, the present agent-based model has allowed us to understand how a set of collaborating agents can better explore the knowledge space in which they are located. Our model, at the same time, has shed light on two important aspects of R&D networks, that can possibly be extended to other collaboration networks: the optimization of the network performance in terms of knowledge exploration, and the emergence of clusters in the knowledge space where the agents interact. Our results, combined with the empirical observation of different alliance formation rates in different industrial sectors [36, 41], could be considered as the first step towards the empirical validation of the performance of R&D alliance networks.

However, more ingredients can be added to the model in order to capture further effects observed on real R&D networks or other kinds of collaboration networks. The first possible extension is the inclusion of more complex strategic link formation rules between the agents, together with the relaxation of the monogamous network approximation, similarly to previous works [26, 27, 42]. In this way, one could investigate network topologies that are closer to the

empirical observations. The alliance formation rules might also be extended with the addition of an inverted U-shaped – rather than linear – relationship between the success of a link activation and the knowledge distance of the two involved partners.

Further extensions consist in the addition of a stochastic term to the agents' motion in the knowledge space, to model the firm self-innovation dynamics, or by the adoption of an open-ended knowledge space, that could be more realistic in high-technology industrial sectors. Another extension is represented by the study of different indicators of the network performance; for instance, one could analyze the share of the knowledge space that has actually been explored by the agents, as opposed to the distance traveled. A more conceptual research question is whether any measure of traveled distance is a better performance indicator than the number of emerging knowledge attractors, which represents instead a more static, equilibrium measure. One could probably take both kinds of indicators into account, or combine them in an appropriate manner; this would surely require a case-by-case discussion, depending on the system under examination.

However, provided that appropriate methodologies are known to locate the interacting agents in a metric knowledge space, this model paves the way for further empirical studies not only on R&D networks, but also on collaboration networks in general. The scope would be to measure knowledge positions and trajectories of agents in real knowledge spaces, using – just to name two prominent examples – patent data for firms, or publication data for scientific authors. In the case of empirical R&D networks, alliance formation rates and knowledge regimes characterizing a set of industrial sectors could be quantified and compared, allowing for a check of the consistency of our model with the observed variations in alliance activities across sectors.

## Acknowledgements

The authors wish to acknowledge the anonymous reviewers, whose comments have greatly improved the quality of this study, in particular the interpretation of our findings. M. V. T. and F. S. acknowledge financial support from the Swiss National Science Foundation (SNF) through grant 100014\_126865, “R&D Network Life Cycles”, and from the ETH Zurich Risk Center Seed Project SP-RC 01-15, “Performance and resilience of collaboration networks”.

## References

- [1] Ahuja, G., Collaboration networks, structural holes, and innovation: A longitudinal study, *Administrative science quarterly* **45** (2000) 425–455.
- [2] Ahuja, G., The duality of collaboration: Inducements and opportunities in the formation of interfirm linkages, *Strategic management journal* **21** (2000) 317–343.
- [3] Axelrod, R., The dissemination of culture, *Journal of conflict resolution* **41** (1997) 203–226.

- [4] Baum, J., Cowan, R., and Jonard, N., Network-independent partner selection and the evolution of innovation networks, *Management Science* **56** (2010) 2094–2110.
- [5] Baum, J. A., Calabrese, T., and Silverman, B. S., Don’t go it alone: Alliance network composition and startups’ performance in Canadian biotechnology, *Strategic management journal* **21** (2000) 267–294.
- [6] Cohen, W. and Levinthal, D., Absorptive capacity: a new perspective on learning and innovation, *Administrative science quarterly* **35** (1990) 128–152.
- [7] Cohen, W. M. and Levinthal, D. A., Innovation and Learning: The Two Faces of R & D, *The Economic Journal* **99** (1989) 569–596.
- [8] Cowan, R., Jonard, N., and Zimmermann, J., Bilateral Collaboration and the Emergence of Innovation Networks, *Management Science* **53** (2007) 1051–1067.
- [9] Das, T. and Teng, B., A resource-based theory of strategic alliances, *Journal of management* **26** (2000) 31.
- [10] Deffuant, G., Neau, D., Amblard, F., and Weisbuch, G., Mixing beliefs among interacting agents, *Advances in Complex Systems* **3** (2000) 87–98.
- [11] DeGroot, M. H., Reaching a consensus, *Journal of the American Statistical Association* **69** (1974) 118–121.
- [12] Fagiolo, G. and Dosi, G., Exploitation, exploration and innovation in a model of endogenous growth with locally interacting agents, *Structural Change and Economic Dynamics* **14** (2003) 237–273.
- [13] Fischer, M. M. and Fröhlich, J., *Knowledge, complexity and innovation systems* (Springer Science & Business Media, 2001).
- [14] Garas, A., Tomasello, M. V., and Schweitzer, F., Selection rules in alliance formation: strategic decisions or abundance of choice? (2014), *ArXiv preprint*, arXiv:1403.3298.
- [15] Gersbach, H. and Schmutzler, A., Endogenous spillovers and incentives to innovate, *Economic Theory* **21** (2003) 59–79.
- [16] Gilbert, N., Agent-based social simulation: dealing with complexity, Technical report, Center for Research on Social Simulation, University of Surrey, Guildford, UK (2004).
- [17] Gomes-Casseres, B., Hagedoorn, J., and Jaffe, A., Do alliances promote knowledge flows?, *Journal of Financial Economics* **80** (2006) 5–33.
- [18] Grant, R. and Baden-Fuller, C., A knowledge accessing theory of strategic alliances, *Journal of Management Studies* **41** (2004) 61–84.
- [19] Groeber, P., Schweitzer, F., and Press, K., How Groups Can Foster Consensus: The Case of Local Cultures, *Journal of Artificial Societies and Social Simulation* **12** (2009) 4.

- [20] Gulati, R., Sytch, M., and Tatarynowicz, A., The Rise and Fall of Small Worlds: Exploring the Dynamics of Social Structure, *Organization Science* **23** (2012) 449–471.
- [21] Hagedoorn, J., Inter-firm R&D partnerships: an overview of major trends and patterns since 1960, *Research policy* **31** (2002) 477–492.
- [22] Hagedoorn, J., Link, A. N., and Vonortas, N. S., Research partnerships, *Research Policy* **29** (2000) 567–586.
- [23] Hanaki, N., Nakajima, R., and Ogura, Y., The dynamics of R&D network in the IT industry, *Research policy* **39** (2010) 386–399.
- [24] Hegselmann, R. and Krause, U., Opinion dynamics and bounded confidence: models, analysis and simulation, *Journal of Artificial Societies and Social Simulation* **5** (2002).
- [25] König, M. D., Battiston, S., Napoletano, M., and Schweitzer, F., The efficiency and stability of R&D networks, *Games and Economic Behavior* **75** (2012) 694–713.
- [26] König, M. D. and Tessone, C. J., Network evolution based on centrality, *Physical Review E* **84** (2011) 056108.
- [27] König, M. D., Tessone, C. J., and Zenou, Y., Nestedness in Networks: A Theoretical Model and Some Applications, *Theoretical Economics* **9** (2014) 695–752.
- [28] Lane, P. J. and Lubatkin, M., Relative absorptive capacity and interorganizational learning, *Strategic management journal* **19** (1998) 461–477.
- [29] Liebeskind, J. P., Knowledge, Strategy, and the Theory of the Firm, *Strategic Management Journal* **17** (1996) 93–109.
- [30] Mowery, D., Oxley, J., and Silverman, B., Technological overlap and interfirm cooperation: implications for the resource-based view of the firm, *Research Policy* **27** (1998) 507–523.
- [31] Nelson, R. R. and Winter, S. G., *An evolutionary theory of economic change* (Cambridge, Mass. : Belknap Press of Harvard University Press, 1982).
- [32] Owen-Smith, J. and Powell, W. W., Knowledge networks as channels and conduits: The effects of spillovers in the Boston biotechnology community, *Organization science* **15** (2004) 5–21.
- [33] Powell, W., White, D., Koput, K., and Owen-Smith, J., Network Dynamics and Field Evolution: The Growth of Interorganizational Collaboration in the Life Sciences, *American Journal of Sociology* **110** (2005) 1132–1205.
- [34] Pyka, A. and Fagiolo, G., *Agent-based modelling: a methodology for neo-Schumpeterian economics* (Elgar companion to neo-schumpeterian economics, 2007).
- [35] Rosenkopf, L. and Almeida, P., Overcoming local search through alliances and mobility, *Management science* **49** (2003) 751–766.

- [36] Rosenkopf, L. and Schilling, M., Comparing alliance network structure across industries: observations and explanations, *Strategic Entrepreneurship Journal* **1** (2007) 191–209.
- [37] Sampson, R. C., R&D Alliances and Firm Performance: the Impact of Technological Diversity and Alliance Organization on Innovation, *Academy of Management Journal* **50** (2007) 364–386.
- [38] Schweitzer, F. and Behera, L., Nonlinear voter models: the transition from invasion to coexistence, *The European Physical Journal B-Condensed Matter and Complex Systems* **67** (2009) 301–318.
- [39] Suzumura, K., Cooperative and Noncooperative R&D in an Oligopoly with Spillovers, *The American Economic Review* **82** (1992) 1307–1320.
- [40] Tessone, C. J. and Zanette, D. H., Synchronised firing induced by network dynamics in excitable systems, *EPL (Europhysics Letters)* **99** (2012) 68006.
- [41] Tomasello, M. V., Napoletano, M., Garas, A., and Schweitzer, F., The Rise and Fall of R&D Networks. (2013), *ArXiv preprint*, arXiv:1304.3623.
- [42] Tomasello, M. V., Perra, N., Tessone, C. J., Karsai, M., and Schweitzer, F., The Role of Endogenous and Exogenous Mechanisms in the Formation of R&D Networks, *Scientific Reports* **4** (2014) 5679.
- [43] Tomasello, M. V., Tessone, C. J., and Schweitzer, F., The effect of R&D collaborations on firms’ technological positions, *International Forum on Knowledge Asset Dynamics (IFKAD) 2015 Proceedings* (2015) 260–276.
- [44] Vazquez, F. and Zanette, D., Epidemics and chaotic synchronization in recombining monogamous populations, *Physica D: Nonlinear Phenomena* **239** (2010) 1922–1928.
- [45] Winter, S. G., Schumpeterian competition in alternative technological regimes, *Journal of Economic Behavior and Organization* **5** (1984) 287–320.